Descriptor Ensemble: An Unsupervised Approach to Descriptor Fusion in the Homography Space

Yuan-Ting Hu · Yen-Yu Lin · Hsin-Yi Chen · Kuang-Jui Hsu · Bing-Yu Chen

arXiv:1412.4196v1 [cs.CV] 13 Dec 2014

Abstract With the aim to improve the performance of feature matching, we present an unsupervised approach to fuse various local descriptors in the space of homographies. Inspired by the observation that the homographies of correct feature correspondences vary smoothly along the spatial domain, our approach stands on the unsupervised nature of feature matching, and can select a good descriptor for matching each feature point. Specifically, the *homography space* serves as the common domain, in which a correspondence obtained by any descriptor is considered as a point, for integrating various heterogeneous descriptors. Both geometric coherence and spatial continuity among correspondences are considered via computing their *geodesic distances* in the space. In this way, mutual verification across different descriptors is allowed, and correct correspondences will be highlighted with a high degree of consistency (i.e., short geodesic distances here). It follows that *one-class SVM* can be applied to identifying these correct correspondences, and boosts the performance of feature matching. The proposed approach is comprehensively compared with the state-of-the-art approaches, and evaluated on four benchmarks of image matching. The promising results manifest its effectiveness.

Keywords Image feature matching · Descriptor fusion · Geometric verification · Homography space · One-class SVM

Yuan-Ting Hu · Yen-Yu Lin · Kuang-Jui Hsu
Research Center for Information Technology Innovation, Academia Sinica, Taipei, Taiwan
E-mail: r01922042@ntu.edu.tw, yylin@citi.sinica.edu.tw, and kjhsu@citi.sinica.edu.tw.

Hsin-Yi Chen · Bing-Yu Chen
Department of Computer Science and Information Engineering, National Taiwan University, Taipei, Taiwan
E-mail: fensi@cmlab.csie.ntu.edu.tw and robin@ntu.edu.tw.

1 Introduction

Image matching aims to identify common regions across images. As a key component of image content analysis, image matching has attracted great attention for several years. It is one of the critical stages in widespread image processing and computer vision applications, such as panoramic stitching (Szeliski and Shum, 1997), object recognition (Lowe, 2004), image retrieval (Datta et al., 2008), and common pattern discovery (Liu and Yan, 2010).

Coupling interest points with local feature descriptors has been proven to be an effective way of image matching (Mikolajczyk and Schmid, 2005; Mikolajczyk et al., 2005). Although the development of powerful descriptors (e.g., Bay et al., 2006; Belongie et al., 2002; Berg and Malik, 2001; Hauagge and Snavely, 2012; Lowe, 2004; Tola et al., 2010; Wang et al., 2011) has gained significant progress, there is in general no a single descriptor that is sufficient for dealing with all kinds of variations and deformations in feature matching. Most descriptors are designed on the trade-off between *distinctiveness* and *invariance*. The more distinctive the descriptor is, the higher precision but the lower recall it may get. On the contrary, descriptors with high degrees of invariance often result in high recall but low precision. It implies that the goodness of a descriptor is usually *image-dependent*. Without any prior knowledge about images, using only one descriptor becomes insufficient and unreliable to conquer the wild image matching problems.

Fig. 1 shows the matching results on two image pairs, `magic cube` and `car`, by using three descriptors, *SIFT* (Lowe, 2004), *raw intensities*, and *geometric blur* (Berg and Malik, 2001), and our approach, respectively. The strong color coherence presents in the case of `magic cube`, so the color-based descriptor, raw in-

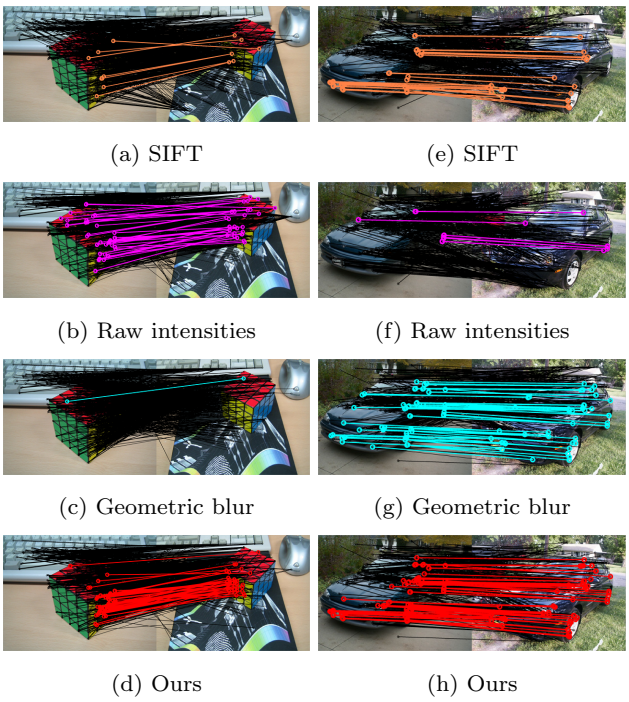

(a) SIFT (e) SIFT

(b) Raw intensities (f) Raw intensities

(c) Geometric blur (g) Geometric blur

(d) Ours (h) Ours

Fig. 1 Feature matching on two image pairs, (a) ∼ (d) `magic cube` and (e) ∼ (h) `car`. The matching results by using three different descriptors, including SIFT, raw intensities, and geometric blur, are shown in the first three rows, respectively. While correct correspondences are drawn in a specific color, wrong ones are in black. In `magic cube`, color/intensity information is important for matching owing to the high degree of color coherence. In contrast, shape and gradient features are more reliable in `car`. This example indicates that the performance of a descriptor varies from image to image. In addition, the deficiency of using a single descriptor is revealed. Our approach instead makes use of multiple, complementary descriptors, and can achieve superior matching results, as shown in (d) and (h).

tensities, gives good results. On the other hand, better performance is achieved by the shape-based descriptors, SIFT and geometric blur, in the case of `car`. However, none of the three descriptors perform well in both the two cases. This example points out not only the performance fluctuation of a descriptor among images but also the deficiency of using a single descriptor.

In view of these issues, we aim at improving the quality of image matching with the aid of multiple, complementary descriptors. Two challenges arise in this scenario. First, features extracted by different descriptors are usually of different dimensions and with different scales of statistics. Even their adopted metrics for similarity measure are diverse. How to effectively fuse heterogeneous descriptors becomes a challenging problem. Second, image matching in general is an unsupervised task. The goodness of descriptors is hard to evaluate without ground truth. When feature matchings by different descriptors present, how to identify correct ones from them is another problem. In this paper, we present an unsupervised approach that can effectively overcome the two problems, and generate accurate and dense correspondences by leveraging complementary descriptors.

The idea of our approach is illustrated through Fig. 2. Fig. 2(a) and 2(b) show two input images for matching. The matching results by three descriptors, including SIFT, raw intensities, and geometric blur, are plotted in Fig. 2(c), 2(d), and 2(e), respectively. Wrong correspondences are drawn in black. Each correct correspondence is displayed with a specific color according to the common object that it resides in. Despite the varieties, a correspondence by any descriptor can be specified by its *geometric transformation* (or *homography*) in the same way. It implies that correspondences by all descriptors can be treated as points in the homography space, which can then serve as the common domain for fusing heterogeneous descriptors. We compute the pair-wise distances among points (correspondences), and show them in Fig. 2(f) via *multi-dimensional scaling* (MDS) (Cox and Cox, 1994), which summarizes high-dimensional data in a low-dimensional space by taking the pair-wise distances as input. The circle, square, and diamond markers denote correspondences obtained by using SIFT, raw intensities, and geometric blur, respectively. It can be observed in Fig. 2(f) that correct (colored) correspondences on the same object share similar homographies no matter by which descriptors they are established. They hence gather together in the homography space, while incorrect correspondences distribute irregularly. This observation suggests that geometric consistency among correspondences is highly relevant to their correctness. Moreover, Fig. 2(f) also indicates that correct correspondences are spatially correlated, since only correct correspondences within the same objects are geometrically consistent. The details of the adopted homography space and the similar measure between correspondences will be introduced later.

Inspired by the observation in Fig.2, this work carries out feature matching with multiple descriptors, and can distinguish itself with the following three main contributions. First, we propose to use the *homography space* as a unified domain, for fusing photometric and geometric information caught by descriptors, where correspondences obtained by all descriptors are represented as points. In this way, the unified representation of correspondences enables geometric checking across diverse descriptors. To our knowledge, this is novel in the field. Second, we present an approach that stands on the unsupervised nature of image matching. It can determine the correctness of correspondences, and choose an appropriate descriptor for matching each feature

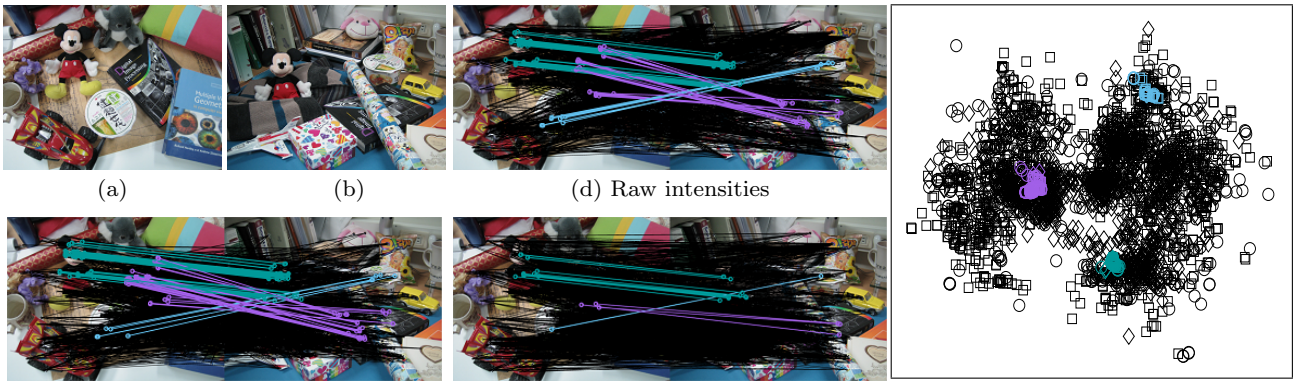

(a) (b) (d) Raw intensities

(c) SIFT (e) Geometric blur (f)

Fig. 2 (a) & (b) Two input images for feature matching. (c) Matching results by using SIFT. (d) Matching results by using raw intensities. (e) Matching results by using geometric blur. Each wrong correspondence is drawn in black, while each correct one is in a specifical color depending on the common object that it resides in. (f) The 2D visualization of correspondences in the homography space via multi-dimensional scaling (MDS). The circle, square, and diamond markers denote correspondences obtained by SIFT, raw intensities, and geometric blur, respectively. This figure demonstrates that not only geometric coherence but also spatial continuity are highly relevant the correctness of correspondences.

point without any training data. Specifically, we estimate the geometric and spatial consistency among correspondences via computing *geodesic distances* (Tenenbaum et al., 2000) on a designed graph to smoothly transfer the carried information in the homography space. Through this process, correct correspondences are highlighted with strong coherence with each other. It follows that *one-class SVM* (Schölkopf et al., 2001) can be applied to picking these correct correspondences. Third, our approach is comprehensively evaluated on four benchmarks of image matching, and jointly takes five descriptors into account, including *SIFT* (Lowe, 2004), *LIOP* (Wang et al., 2011), *DAISY* (Tola et al., 2010), *geometric blur* (GB) (Berg and Malik, 2001) and *raw intensities* (RI). Our approach is compared with four image matching baselines and four baselines of descriptor fusion. It achieves significantly better results than those by the best descriptors and baselines in most cases. We hence term our approach *Descriptor Ensemble* in the sense it combines multiple, complementary descriptors, and improves the performance of image matching.

The rest of this manuscript is organized as follows. A review of the related works is given in Section 2. The problem we address is stated in Section 3. We described the adopted homograph space and the used similarity measure between correspondences in Section 4. The proposed approach is specified in Section 5. The experimental setup and results are given in Section 6 and Section 7, respectively. Finally, we conclude this work in Section 8.

2 Related Work

The literature on image feature matching is quite extensive. Our review focuses on those crucial to the development of the proposed approach.

2.1 Local Feature Descriptors

Local feature descriptors (Mikolajczyk and Schmid, 2005) have been extensively studied, especially since the seminal works by Schmid and Mohr (1997) and Lowe (2004). Various descriptors have been designed to be robust to noises while invariant to particular types of deformations in matching. For example, *SIFT (scale-invariant feature transform)* (Lowe, 2004) describes image regions in the gradient domain, constructing an 128-dimensional histogram, and is known to be robust to scale and orientation changes. *LIOP (local intensity order pattern)* (Wang et al., 2011) encodes both the local and global ordinal information, and can alleviate the unfavorable variations caused by the changes of lighting conditions. *DAISY* (Tola et al., 2010) is featured with fast feature extraction, while keeps invariant to viewpoint changes. In addition, diverse visual cues have been explored in descriptor construction, such as color characteristics (Abdel-Hakim and Farag, 2006), shapes (Berg and Malik, 2001), internal self-similarities (Shechtman and Irani, 2007), topological information (Lobaton et al., 2011), and local symmetries (Hauagge and Snavely, 2012). These descriptors are designed on the trade-off between distinctiveness and invariance. Thus, there does not exist an optimal descriptor in a wide range of test images. By contrast,

we introduce our approach into image matching by employing multiple descriptors to complement one another and, thus solve this problem.

Instead of using handcrafted descriptors, a branch of research efforts has been made to automatically derive discriminative descriptors by machine learning techniques (e.g., Carneiro, 2010; Hua et al., 2007; Varma and Ray, 2007; Winder and Brown, 2007). Hua et al. (2007) introduced discriminant learning into constructing feature descriptors by minimizing the distances of correct matching pairs and maximizing the distances of non-matched ones. Varma and Ray (2007) found an optimal tradeoff for a given classification problem and proposed a kernel learning method of learning transformations of features from training image pairs. Carneiro (2010) proposed a *universal feature transform* that maps patches into robust descriptors. The universal feature transform can be incrementally updated, and hence expand its applicability. However, additional training data are required in these approaches. Moreover, since the optimal descriptors for matching vary from image to image, the performances of these learned descriptors highly rely on the consistency between training images and test images. Instead, our approach provides a fully unsupervised way to select proper descriptors. Neither additional training data nor prior knowledge about the images are needed.

2.2 Correspondence Verification

Identifying correct feature correspondences from candidates is an important step in image matching. Geometric layout checking is one of the most effective ways, because the geometric layout of feature correspondences often reveals their correctness. *RANSAC* (Fischler and Bolles, 1981) is a classic method for removing outliers through geometric verification. It estimates a global transformation and rejects outliers simultaneously. A correspondence is considered as an outlier and deleted if it is inconsistent with the transformation that the majority agree. One advantage of RANSAC is its easy implementation to fulfill geometric verification. However, RANSAC is not able to deal with multiple object matching and non-rigid transformations, and would be computationally expensive when the number of outliers becomes large.

Chui and Rangarajan (2003); Ma et al. (2014) and Pang et al. (2014) relaxed the geometric assumption of correspondences from obeying a global transformation to a smooth feature mapping function. The feature points are linked to their corresponding points through a smooth mapping function, which makes a non-rigid transformation expressible. Ma et al. (2014) and Pang et al. (2014) have demonstrated an effective way to determine the parameter values of the mapping function with the *vector field*. They can identify the correct correspondences. However, owing to the smoothness assumption of the mapping function, these methods are not very robust in matching multiple objects, especially when the transformations of multiple objects vary greatly.

Instead of deriving the transformations involved in matching, non-rigid deformations can be dealt with via measuring the similarities between correspondences, since correct correspondences tend to be consistent with each other. Both the photometric information given by descriptors and the geometric relationship of correspondences can be used in similarity computation. Some examples of similarity measures can be found in works by Choi and Kweon (2009); Leordeanu and Hebert (2005); Zheng and Doermann (2006) and Liu and Yan (2010). With the similarities between correspondences, a branch of research efforts (Cho et al., 2010; Choi and Kweon, 2009; Cour et al., 2006; Leordeanu et al., 2012; Li et al., 2013; Zheng and Doermann, 2006) cast the task of correct correspondence identification as an optimization problem over the matching score. *Graph matching* (Cho et al., 2010; Cour et al., 2006; Leordeanu et al., 2012) is one of the representative techniques in optimizing the matching score. Although it is an NP hard problem, various graph matching algorithms have been proposed to get the approximate solutions. Nonetheless, these approaches are sensitive to outliers, and are less robust in multiple object matching.

Clustering based techniques for grouping correspondences with geometric constraints have been explored. Cho et al. (2009) established a linkage model of correspondence clusters, and iteratively merged the clusters based on their geometric consistency. Zhang et al. (2012) refined and reformulated the linkage model as a directed graph to further eliminate ambiguousness. The computational efficiency might be an issue due to the iterative algorithms. Liu and Yan (2010) found local maximizers on the matching score and merged them if any two of them are similar enough. The clustering framework can handle multiple transformations, but the parameter values of the developed models are difficult to determine, such as the thresholds of identifying the correct clusters, the criteria of merging clusters and the scale factor in (Liu and Yan, 2010).

Voting schemes can be used for measuring the consistency between correspondences. The correspondences are checked by the pair-wise geometric consistency via mutual voting among correspondences. Avrithis and Tolias (2014) transformed correspondences into Hough space and the voting results

are collected efficiently with a pyramid structure. Chen et al. (2013) cast the voting process as a kernel density estimation problem in the transformation space. These approaches can identify multiple objects effectively. However, voting methods would become less powerful to find correct correspondences when the number of correct matchings is so less that votes from them are not dominant during voting. Our proposed approach in spirit follows the idea of project correspondences into transformation space, but we use the geodesic distance to include the spatial layout for improving the geometric consistency measurement. The most important feature of our approach is that multiple descriptors are considered. The proposed approach increases the number of correct matchings, and can avoid the situation mentioned above.

2.3 Multiple Descriptor Fusion

Since different descriptors can catch diverse visual cues, using multiple descriptors has been a feasible way for improving performance. A number of approaches (such as Bosch et al., 2007; Brox and Malik, 2011; Harada et al., 2010; Mortensen et al., 2005; Weinzaepfel et al., 2013; Zhang et al., 2007) have been developed to fuse diverse descriptors for improving image matching, classification and alignment. Mortensen et al. (2005) proposed to *concatenate* SIFT (Lowe, 2004) and shape context (Belongie et al., 2002), and reported good results in matching. However, simple feature concatenation ignores the possible variations of feature dimensions and feature value scales among descriptors. It may lead to suboptimal performance, especially when less powerful descriptors have dominant feature dimensions or values. To address this problem, Bosch et al. (2007) represented images under each descriptor as a *kernel matrix*. The works by Brox and Malik (2011) and Weinzaepfel et al. (2013) integrated descriptor matching for handling large displacement in optical flow, and combined multiple visual evidences represented in form of *energy functions*. Although kernel matrices and energy functions can serve as the unified domains for descriptor fusion, these approaches tune or learn fixed weights for descriptor combination. It may not be suitable for image matching, because the optimal descriptors change from image to image. Furthermore, using brute force search to determine descriptor weights may become infeasible, when there are a large number of descriptors to be considered. Besides, image matching is an unsupervised task, and no training or validation data are available for determining descriptor weights in general.

Our approach tackles these issues, and has the following two advantages: 1) Multiple descriptors are represented by their homographies in matching so that we can work with complementary descriptors without worrying about their diversities of feature dimensions or feature value scales; 2) Our approach allows geometric checking across descriptors, and consensus correspondences will reveal through the process. It means that the plausible correspondences by various descriptors can be jointly identified in a fully unsupervised manner.

3 Problem Statement

We aim to match two given images I^P and I^Q, which come with the sets of detected feature points, $U^P = \{u_i^P\}_{i=1}^{N^P}$ and $U^Q = \{u_i^Q\}_{i=1}^{N^Q}$, respectively. The support region of each feature $u_i \in U^P \cup U^Q$ is assumed to an ellipse in this work. These feature points can be obtained by using off-the-shelf detectors, such as *Harris-Affine* (Mikolajczyk and Schmid, 2004), *Hessian-Affine* (Mikolajczyk and Schmid, 2004), the salient region detector (Kadir et al., 2004), or their combinations. We use Hessian-Affine detector for its efficiency and high repeatability. Multiple descriptors are employed to characterize each feature point. The center and the described appearances of feature u_i are respectively denoted by $\mathbf{x}_i$ and $\{\mathbf{f}_{i,m}\}_{m=1}^{M}$, where M is the number of the employed descriptors. The set $\tilde{\mathcal{C}} = U^P \times U^Q$ covers all possible feature *correspondences* (or *matchings*). Our goal is to detect correct correspondences in $\tilde{\mathcal{C}}$ as many as possible.

The number of the detected feature points in a image, i.e., N^P or N^Q, is often in the order of 10^3. Directly searching correct correspondences in $\tilde{\mathcal{C}}$ may be inefficient. Hence we start from compiling a reduced set $\mathcal{C}$ of $\tilde{\mathcal{C}}$ by filtering out correspondences that are unlikely to be correct. For each feature $u_i^P \in I^P$, we find the set of the most similar r matchings, $\mathcal{C}_{i,m} = \{(u_i^P, u_{i_k}^Q \in I^Q)\}_{k=1}^{r}$, with descriptor m and the yielded distance $\|\mathbf{f}_{i,m}^P - \mathbf{f}_{i_k,m}^Q\|$. Since total M descriptors are adopted, at most $r \times M$ matchings of u_i^P are kept in $\mathcal{C}$ after removing the duplicates. Namely,

$$\mathcal{C} = \bigcup_{i=1}^{N^P} \mathcal{C}_i, \text{ where } \mathcal{C}_i = \bigcup_{m=1}^{M} \mathcal{C}_{i,m}. \tag{1}$$

We will work on the reduced candidate set $\mathcal{C}$. The value of r controls the trade-off between efficiency and accuracy. It is empirically set as 5 in all our experiments.

4 Homography Space: A Common Domain for Descriptor Fusion

In this section, we first introduce how to compute the geometric transformations of correspondences and how to measure their geometric dissimilarity. Then, we show how to fuse diverse descriptors in the homography space, where correspondences obtained by different descriptors are treated in a unified manner.

4.1 Homographies of Feature Correspondences

The elliptical region of feature u_i can be specified by mapping a circular region centered on the origin via the affine transformation:

$$T(u_i) = \begin{bmatrix} A(u_i) & \mathbf{x}_i \\ \mathbf{0}^\top & 1 \end{bmatrix} \in \mathbb{R}^{3\times 3}, \tag{2}$$

where $\mathbf{x}_i \in \mathbb{R}^{2\times 1}$ is the feature center, and $A(u_i) \in \mathbb{R}^{2\times 2}$ is a non-singular matrix which accounts for the scale, the shape, and the orientation of u_i. After normalization with transformation $T(u_i)^{-1}$, all the adopted descriptors can be applied to u_i, and generate $\{\mathbf{f}_{i,m}\}_{m=1}^{M}$. Refer to work by Mikolajczyk and Schmid (2004) for the details.

A homography in this work refers to the geometric transformation of a feature correspondence. For a correspondence between $u_i^P \in U^P$ and $u_j^Q \in U^Q$, the transformation from the support region of u_i^P to that of u_j^Q can be derived as

$$H_{ij} = T(u_j^Q) * T(u_i^P)^{-1} \in \mathbb{R}^{3\times 3}. \tag{3}$$

H_{ij} is a 6-dof (degrees of freedom) affine transformation. Thus, it can be viewed as a point in the 6-dimensional homography space $\mathcal{H}$.

Consider two correspondences $c = (u_i^P, u_j^Q) \in \mathcal{C}$ and $c' = (u_{i'}^P, u_{j'}^Q) \in \mathcal{C}$. We use the *reprojection error* to measure their geometric dissimilarity. Specifically, the homography matrices, H_{ij} and $H_{i'j'}$, of the two correspondences are firstly computed by Eq. (3). The *projection error* of (u_i^P, u_j^Q) with respect to $H_{i'j'}$ is then calculated by

$$d_{err}(u_i^P, u_j^Q, H_{i'j'}) = \|\mathbf{x}_j^Q - \rho(H_{i'j'} \begin{bmatrix} \mathbf{x}_i^P \\ 1 \end{bmatrix})\|, \tag{4}$$

where function $\rho : \mathbb{R}^3 \to \mathbb{R}^2$ is defined as

$$\rho(\begin{bmatrix} a \\ b \\ c \end{bmatrix}) = \begin{bmatrix} \frac{a}{c} \\ \frac{b}{c} \end{bmatrix}. \tag{5}$$

The projection error $d_{err}(u_i^P, u_j^Q, H_{i'j'})$ is the induced error when changing the homography from H_{ij} to $H_{i'j'}$ on correspondence (u_i^P, u_j^Q). The reprojection error between correspondences c and c' is then defined as

$$\begin{aligned} d(c, c') = &\frac{1}{4}\big(d_{err}(u_i^P, u_j^Q, H_{i'j'}) + d_{err}(u_j^Q, u_i^P, H_{i'j'}^{-1}) \\ &+ d_{err}(u_{i'}^P, u_{j'}^Q, H_{ij}) + d_{err}(u_{j'}^Q, u_{i'}^P, H_{ij}^{-1})\big). \end{aligned} \tag{6}$$

We will use the reprojection error to measure the geometric dissimilarity between correspondences in $\mathcal{C}$.

4.2 Homography as A Unified Representation

In view of the complexity of feature matching, we consider characterizing each feature point u_i with total M kinds of different descriptors, i.e., $\{\mathbf{f}_{i,m} \in \mathcal{F}_m\}_{m=1}^{M}$, and each descriptor is associated with a measure of photometric dissimilarity $d_m : \mathcal{F}_m \times \mathcal{F}_m \to \mathbb{R}$.

The resulting representations by these descriptors are typically of various dimensions and even in diverse forms, such as vectors, histograms, and pyramids. The associated dissimilarity measures may be also different, such as Euclidean distance, χ^2 distance or negative cosine similarity. To avoid the difficulties caused by these varieties, a unified representation of these descriptors is preferred.

It can be observed that the homography of a feature correspondence in Eq. (3) is *descriptor-independent*, and can hence serve as the domain for descriptor fusion. That is, each descriptor determines its own candidate correspondences as shown in Eq. (1), while all the candidate correspondences are represented by the corresponding homographies, each of which can be treated as a point in the homography space $\mathcal{H}$ of six dimensions. In this way, the dissimilarity between correspondences that are generated by different descriptors can be measured through Eq. (6). It implies that knowledge, such as the geometric distribution of matchings, can be transferred across descriptors. The proposed approach will leverage this nice property to work with complementary descriptors, and boost the performance of image matching.

5 The Proposed Approach

We formulate the task of image matching as finding an appropriate matching for each feature point u_i^P in image I^P, i.e., picking the most plausible correspondence from $\mathcal{C}_i$ in Eq. (1). In this work, multiple descriptors collaborate in the sense that they jointly determine candidate correspondences with diverse region characteristics and different kinds of invariance, so the probability that the correct correspondence resides in

$\mathcal{C}_i$ largely increases. The goal at this stage is to determine the correct correspondence for each u_i^P, if it exists. The unsupervised nature of image matching makes this task very challenging, because no prior knowledge or relevant training data are available.

We tackle this issue based on the observation that the homographies of correct correspondences vary smoothly along the spatial locations in the image. Specifically, we firstly employ a graph to encode the spatial arrangement among correspondences, and compute the *geodesic distances* (Tenenbaum et al., 2000) on the graph for geometric coherence estimation. Then, we utilize *one-class SVM* (Schölkopf et al., 2001) to identify correct correspondences, since it can effectively separates alike (both geometrically and spatially coherent here) data from the outliers. The two steps are respectively described in the following.

5.1 Geometric and Spatial Coherence Estimation

To jointly consider the geometric and spatial relationships among correspondences, we construct a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, in which each correspondence $c_i \in \mathcal{C}$ is associated with a vertex $v_i \in \mathcal{V}$, while an undirected edge $e_{ij} \in \mathcal{E}$ is added to connect vertices v_i and v_j if the end points in image I^P of c_i and c_j are near enough. That is, the number of vertices in $\mathcal{G}$, $|\mathcal{V}|$, is the same as the number of correspondences in $\mathcal{C}$. In our implementation, we consider two feature points in I^P are nearby if one point belongs to the k spatial nearest neighbors of the other point. Weight w_{ij} assigned to edge e_{ij} is defined as

$$w_{ij} = \begin{cases} d(c_i, c_j), & \text{if } e_{ij} \in \mathcal{E}, \\ \infty, & \text{otherwise}, \end{cases} \tag{7}$$

where the geometric dissimilarity $d(c_i, c_j)$ between correspondences c_i and c_j is given in Eq. (6). With the weighted graph, we compute the geodesic distance between each pair of vertices (i.e., correspondences). We denote the geodesic distance between correspondences c_i and c_j by $d_{geo}(c_i, c_j)$ hereafter. It can be seen that graph $\mathcal{G}$ integrates the spatial continuity into the estimation of geometric coherence. The use of geodesic distance catches the phenomenon that the homographies of correct correspondences on the same object may change smoothly along their spatial locations. Therefore, the resulting dissimilarity measure can deal with possible deformations in matching.

Suppose there are N correspondence candidates, i.e., $\mathcal{C} = \{c_1, ..., c_N\}$. We compute the pair-wise geodesic distances $\{d_{geo}(c_i, c_j)\}_{i,j=1}^N$. The correct correspondences will be highlighted with strong geometric and spatial coherence (short geodesic distances) with other correct correspondences. It is worth pointing out that compared with $d(c_i, c_j)$ in Eq. (6), the geodesic distance $d_{geo}(c_i, c_j)$ can more faithfully measure the dissimilarity between c_i and c_j, since the spatial information is taken into account to remove the noises, i.e., incorrect correspondences whose homographies happen to be consistent with those of the correct ones. The performance gain of using geodesic distances over reprojection errors will also be evaluated in each of the used dataset in the experiments.

5.2 Correct Matching Identification by One-class SVM

At this stage, we apply one-class SVM to distinguishing the correct correspondences from candidate set $\mathcal{C}$. One-class SVM is one of the state-of-the-art methodologies for unsupervised classification. It separates positive and negative data in an asymmetrical scenario: it assumes that positive data are similar to each other, while negative data are different in their own ways.

In our case, the correct correspondences are usually geometrically and spatially consistent with each other, i.e., short geodesic distances among them. On the other hand, the wrong matchings are caused by various factors, so their homographies often irregularly distribute. It results in that the homography of a wrong correspondence tends to be dissimilar to most homographies of all the other correspondences. Thus, our case closely meets the scenario of one-class SVM: the correct and incorrect matchings respectively correspond to positive and negative data. For the set of correspondence candidates, $\mathcal{C} = \{c_1, ..., c_N\}$, one-class SVM predicts their labels by solving the following constrained optimization problem

$$\begin{aligned} \min_{\mathbf{w}, \{\epsilon_i\}} \quad & \frac{1}{2}||\mathbf{w}||^2 + \frac{1}{Co \cdot N} \sum_{i=1}^{N} \epsilon_i - \nu \\ \text{subject to} \quad & \mathbf{w}^\top \phi(c_i) \geq \nu - \epsilon_i, \text{ for } 1 \leq i \leq N, \\ & \epsilon_i \geq 0, \text{ for } 1 \leq i \leq N, \end{aligned} \tag{8}$$

where Co and ν are the two parameters of one-class SVM. We tuned and fixed them as $Co = 1$ and $\nu = 0.5$ for all the experiments. As a kernel machine, one-class SVM can work on nonlinearly mapped data. The function ϕ maps the data, correspondences here, to a *Reproduced Kernel Hilbert Space* (RKHS), which is implicitly defined by the adopted kernel. The optimization of one-class SVM can be accomplished by referencing only the inner products of pairs of the mapped data, and the inner product can be efficiently computed via the *kernel trick*, i.e., $k(c_i, c_j) = \langle \phi(c_i), \phi(c_j) \rangle$.

Note that we don't explicitly define the feature representation of correspondences in $\mathcal{C}$, but their pair-wise dissimilarity through d_{geo}. A kernel function is used to encode the similarity among data. In this work, the kernel matrix $K \in \mathbb{R}^{N \times N}$ and the kernel function $k(\cdot, \cdot)$ are defined as follows:

$$K(i,j) = k(c_i, c_j) = \exp\left(-\frac{d^2_{geo}(c_i, c_j)}{\sigma^2}\right), \text{ for } 1 \leq i, j \leq N, \quad (9)$$

where σ is the hyperparameter. We set σ as twice of the average geodesic distance from each correspondence to its nearest neighbor. It follows that each correspondence c_i is predicted via $\text{sign}(f(c_i))$, where score $f(c_i)$ is in form of:

$$f(c_i) = \mathbf{w}^\top \phi(c_i) - \nu = \sum_{j=1}^{N} \alpha_j k(c_i, c_j) - \nu, \quad (10)$$

and $\{\alpha_j\}_{j=1}^N$ are the optimized coefficients of the support vectors. Note that the results of image matching are often jointly measured by *precision* and *recall*. For each feature point u_i^P in image I^P, we pick its correspondence as the one that has the highest prediction score in $\mathcal{C}_i$ (cf. Eq. (1)). All picked correspondences are further sorted according to their prediction scores. Precision-recall analysis can then be carried out with the sorted list and a set of thresholds.

Our approach is featured with its high degree of flexibility in matching images with multiple descriptors. First, any elliptical interest region detectors or their combination can be used to locate feature points. Second, any region descriptors can be applied and fused. No assumption about their feature representations, dimensions, the associated similarity measures has been made. Third, motivated by the observation that the optimal descriptor for matching is image-dependent or even region-dependent, our approach allows each feature point to select its correspondence obtained by any descriptor only if the resulting homography is geometrically and spatially consistent. Besides, our approach is easy to implement. The geodesic distances can be computed via *Dijkstra's algorithm* with an approximate version to only update 200 times on each distance for computational efficiency, and a few packages of one-class SVM are available, such as *LibSVM* (Chang and Lin, 2001) which is adopted in this work. The time complexity of our method is between $O(N^2)$ to $O(N^3)$, i.e., the complexity of one-class SVM, where N is the number of correspondences. Note that in this work, N grows linearly with respect to the number of descriptors used for fusion.

6 Experimental Setup

In this section, we introduce the details of our experimental setting, including the used feature detector, descriptors, baselines, datasets, and evaluation criteria.

6.1 Feature Detector and Feature Descriptors

The Hessian affine invariant detector (Mikolajczyk and Schmid, 2004) is used in our experiments to detect interest points and their surrounding elliptical support regions. Each detected region is normalized and rotated to the principal orientation. We apply all the feature descriptors to the normalized patches except for geometric blur. We will explain it later. The average number of detected interest points in an image is around 10^3.

In the experiments, we adopt five feature descriptors, because they capture diverse image characteristics and tend to complement each other. They are listed below in bold and in abbreviation:

SIFT: The *SIFT* (Lowe, 2004) descriptor constructs a 3D histogram on the gradient map, and stacks the 3D histogram as a 128-dimensional vector. This descriptor is known to be rotation and scale invariant.

LIOP: The *LIOP* (Wang et al., 2011) descriptor models the intensity relationship of a patch, and thus is able to handle monotonic intensity changes. The LIOP features are represented by a 144-dimensional vector.

DAISY: The *DAISY* (Tola et al., 2010) descriptor is known for its computational efficiency and invariance to viewpoint changes. It employs spatially varied kernels, and puts high weights on the region near the feature center. The dimension of a DAISY descriptor is 136.

RI: The *Raw intensity* (RI) descriptor is represented by the pixel intensities in a raster scan order. It is the most distinctive descriptor among the five descriptors, but is also the most sensitive to noises and variations. We normalize each support region to 31×31 pixels. Thus, a 961-dimensional RI descriptor is constructed.

GB: The *geometric blur* (GB) descriptor (Berg and Malik, 2001) is widely used for shape matching. Like DAISY, it uses spatially varied kernels to deal with the possible deformation. To make the GB descriptor more complementary to others, we enlarge each detected support region by three times, and apply it to the enlarged region to capture shape information with a wider range.

Euclidean distance is used to measure the distance between two regions under each descriptor.

6.2 Baselines

For performance comparison, we implemented eight baselines of image feature matching. Four of them are the state-of-the-art feature matching algorithms, including *spectral matching* (SM) (Leordeanu and Hebert, 2005), *agglomerative correspondence clustering* (ACC) (Cho et al., 2009), *Hough voting* (HV) (Chen et al., 2013), and *vector field consensus* (VFC) (Ma et al., 2014). The other four are the baselines that fuse multiple descriptors, including *concatenation* (CAT), *concatenation with Hough voting* (CAT+HV), *ranking*, and *ratio*. The following describes the eight baselines.

SM: Leordeanu and Hebert (2005) employed a graph structure to realize geometric checking. Their method relaxes the formulation of an integer quadratic programming problem for feature matching, and obtains an approximate solution by solving an eigenvalue problem. In our implementation, the affinity matrix among correspondences is computed by using a modified RBF function in which Euclidean distance is replaced with the reprojection error.

ACC: Cho et al. (2009) proposed to iteratively group correspondence clusters based on geometric consistency. The reprojection error is used to measure the geometric consistency, as used in (Cho et al., 2009).

HV: Chen et al. (2013) conducted Hough transform for every correspondence and verified its geometric coherence with nearby correspondences in the transformation space. The reprojection error is used in kernel density estimation for geometric verification.

VFC: Ma et al. (2014) proposed to identify outliers by modeling inliers with a vector field. We used the implementation provided by Ma et al. (2014) in the experiments.

CAT: This baseline fuses multiple descriptors by concatenating all the feature descriptors. Before concatenation, the standard deviation of pair-wise distances between feature points under each descriptor is firstly computed, and each descriptor is normalized by dividing by the standard deviation. The nearest neighbor search is applied to finding the possible matching candidates.

CAT+HV: This baseline uses the initial candidates constructed by baseline CAT as the input of HV. Compared with CAT, CAT+HV additionally realizes geometric checking by Hough voting.

Ranking: In this baseline, we compute the nearest neighbors of all feature points in image I^p with a specific descriptor, and assign the ranks to these points according to their distances to the nearest neighbors. The procedure is repeated for each descriptor. For each feature point in I^p, we determine its correspondence based on the descriptor that has the highest rank at this point.

Ratio: For each point in I^p, we find its first two nearest neighbors with a specific descriptor, and compute the distance ratio between the first two neighbors. The smaller the ratio is, the more confident the descriptor is at this point. The procedure is repeatedly done for each descriptor. For this feature point, we determine its correspondence by using the descriptor with the smallest distance ratio.

It is worth mentioning that baselines CAT and Ranking consider all the descriptors are equally important, so their feature vectors or ranks can be directly concatenated or compared. However, the performances of these descriptors in fact vary from image to image. The performance of baseline Ratio degrades when there is high variation among the distance distributions of all the descriptors. The reason is that the distance ratios cannot be compared across different descriptors in that situation.

The first four baselines take one descriptor into account at a time, while the other four baselines and our approach consider all the adopted descriptors jointly. For fair comparison, all the baselines and our approach use the same detected interest regions, descriptors, and evaluation criteria in the experiments.

6.3 Datasets

The performance of our approach is evaluated on four benchmarks of image matching, including Co-recognition (Co-reg for short) dataset (Cho et al., 2008), Object dataset (Cho et al., 2010), Symfeat (SYM for short) dataset (Hauagge and Snavely, 2012) and VGG dataset (Mikolajczyk and Schmid, 2005). To get preeminent performance on all of them is very challenging, because the four datasets contain a variety of variations, such as diverse kinds of deformation, various types of scenes, and different numbers of common objects. They jointly serve as a good test bed for performance evaluation.

Co-reg dataset: It consists of six image pairs. There are multiple common objects in every pair. Large changes in viewpoints, rotations and scales combined with noises coming from the clutter backgrounds and occlusions make matching quite difficult on this dataset.

Object dataset: It gathers 30 pairs of images. Each pair contains different object instances of the same category. The performance of a descriptor fluctuates a lot between different image pairs of this dataset due to the diversity of intra-class variations.

SYM dataset: It is composed of 46 image pairs of architectural scenes. The dramatic variations of this

dataset range from lighting conditions (day/night), ages (old/nowadays scene) to rendering styles (photograph/drawing).

VGG dataset: It contains eight image sets, each of which contains images with five different degrees of a specific type of variation. Total five types of variation are included in the dataset, namely viewpoint changes, scale changes, image blur, JPEG compression, and illumination changes. For each image set, we carry out image matching with two different degrees of variation in the experiment, i.e., the first image vs. the second one and the first one vs. the fourth one, which represent the slight and drastic variations in matching, respectively.

6.4 Evaluation criteria

For performance measure, the evaluation metrics used in the experiments are introduced. For datasets VGG and SYM, we follow Mikolajczyk and Schmid (2005), and consider that a correspondence is correct if the overlap error is less than 50%. For datasets Object and Co-reg where manually annotated ground truth is available, a correspondence is correct if the distance between the matched feature point and the ground truth is within eight pixels.

After determining the correctness of correspondences, the performance of a matching algorithm can be measured by jointly taking *precision* and *recall* into account. While precision is the fraction of detected correspondences that are correct, recall is the fraction of correct correspondences that are detected. Specifically, the two terms are respectively defined as

$$\text{PRECISION} = \frac{n\text{TP}}{n\text{TP} + n\text{FP}}, \tag{11}$$

and

$$\text{RECALL} = \frac{n\text{TP}}{n\text{TP} + n\text{FN}}, \tag{12}$$

where $n\text{TP}$ and $n\text{FP}$ are the numbers of correctly and wrongly detected correspondences by a matching method, respectively. $n\text{FN}$ is the number of correct correspondences that are not detected.

For each matching approach, including ours and the eight adopted baselines, all the detected correspondences are sorted by its own criterion, such as the element values of the eigenvector in baseline SM, the ratio values in baseline Ratio, and the prediction score, Eq. (10), in our approach. By sampling on the sorted lists, the performance of each approach can then be represented by a *precision-recall curve* or *mean average precision* (mAP).

7 Experimental Results

The performance of the proposed approach is evaluated and analyzed in this section. Totally, three sets of experiments are conducted. First, the transformation space is visualized to verify that correct correspondences established by all the descriptors gather together in that space, while incorrect ones distribute irregularly. We also visualize the homography spaces when the reporjection error and the proposed geodesic distance serve as the distance functions, respectively. Second, our approach is compared with the eight baselines on the four benchmarks of feature matching. The obtained quantitative results of all methods are presented in the form of mAPs and precision-recall curves. Third, we show the matching results to demonstrate that our approach can leverage multiple, complementary descriptors to achieve remarkable performance improvement in feature matching.

7.1 Homography Space Visualization

To check whether the homography space is qualified to serve as a uniform domain for descriptor fusion, we visualize it by *multi-dimensional scaling* (MDS) (Cox and Cox, 1994), which can approximate the pair-wise distances between data in a 2D space.

Fig. 3(a) and Fig. 3(b) show the sets of the initial matching candidates, via Eq. (1), by using SIFT and by using all the five descriptors, respectively. There are four common objects in this pair of images. The correct matchings are colored according to the objects that they lie on. By using the geodesic distance presented in Section 5.1, the two sets of matchings are visualized in Fig. 3(c) and Fig. 3(d), respectively. It turns out that no matter how many descriptors are used, the correct (colored) correspondences gather together, while wrong matchings distribute irregularly. This example also points out that using multiple descriptors helps to find out the correct correspondences that are sparsely detected with a single descriptor and may be neglected, such as the purple correspondences in Fig. 3(c). In Fig. 3(d), the purple correspondences given by diverse descriptors can mutually support in both geometric and spatial coherence estimation. It implies that they are more probably predicted as correct correspondences by one-class SVM.

As described in Section 5.1, the developed geodesic distance computed over the designed graph takes both geometric consistency and spatial continuity into account. We compare the geodesic distance with the reprojection error by visualizing the homography spaces that they induce. In Fig. 4(a), two images to be matched

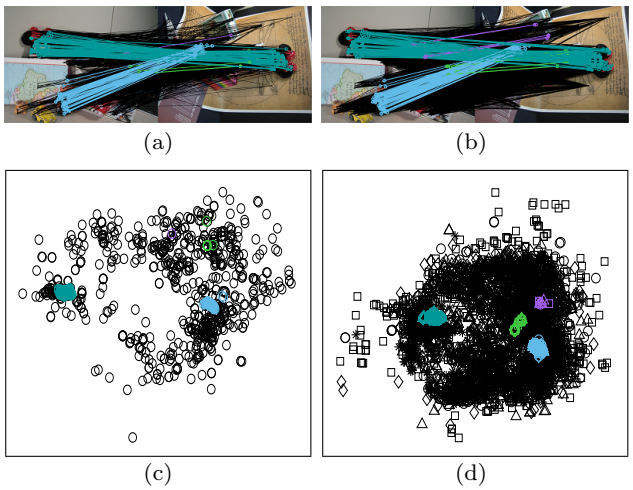

Fig. 3 (a) The matching candidates given by SIFT on image pair `Minnies` of the Co-reg dataset. (b) The matching candidates given by all the five descriptors on the same image pair. (c) 2D visualization of the homograph space of the matchings in (a). (d) 2D visualization of the homograph space of the matchings in (b).

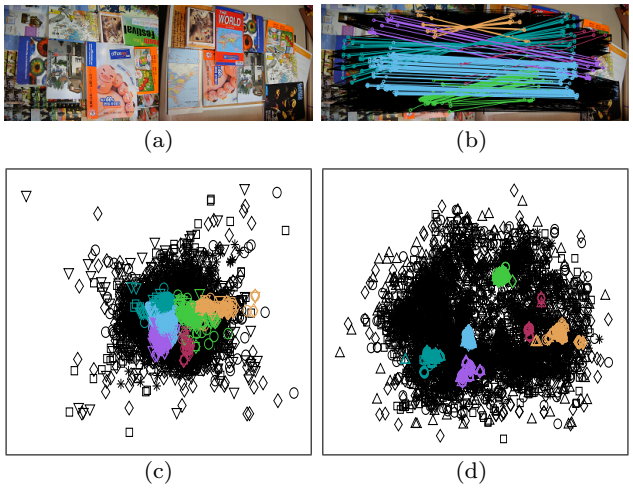

Fig. 4 (a) An image pair, `Books` of the Co-reg dataset. (b) The initial correspondences. (c) 2D visualization of these correspondences in the homography space when the reprojection error is used. (d) 2D visualization of these correspondences in the homography space when the developed geodesic distance is used.

are shown. The initial correspondence candidates are given in Fig. 4(b). We calculate the dissimilarity between these correspondences by using the reprojection error and the geodesic distance, and show their distributions in the homography space in Fig. 4(c) and 4(d), respectively. It can be observed in Fig. 4(c) that the correct correspondences on an object may mix with the correct correspondences on the other objects. On the other hand, the correct correspondences on an object tightly assemble in Fig. 4(d), and are well separated from the rest. It implies that the developed geodesic distance can effectively consider both geometric and spatial consistency to better discover object-aware homographies. This property facilitates correct correspondence identification by one-class SVM.

Table 1 The accuracy in mAP of using the reprojection error and the geodesic distance in our approach

mAP (%)	Co-reg	Object	SYM	VGG
Reprojection Error	73.13	35.97	32.05	88.25
Geodesic Distance	78.46	40.88	37.31	88.39

7.2 Quantitative Results

Following the previous discussion, we evaluate the performance of using the proposed geodesic distance and using the reprojection error in our approach, and report the accuracies in mAP in Table 1. The mAPs on the four datasets are improved when the geodesic distance is used. It indicates that the geodesic distance more faithfully grasps the intrinsic relationships between correspondences by exploring the graph which encodes both the spatial and geometric consistency. The performance gains, about 5%, on the first three datasets, i.e., Co-reg, Object, and SYM, are remarkable. The main reason is that the multiple common objects in Co-reg and the foregrounds and backgrounds in Object and SYM typically have diverse transformations in matching, but each of these transformations tends to vary smoothly in the spatial domain. Hence, modeling spatial coherence is helpful. On the other hand, all interest points in each image of dataset VGG almost undergo the same transformation in matching. Giving additional spatial information does not help much in the cases.

We evaluate and compare our approach, *Descriptor Ensemble* or DE for short, with the eight baselines. Five different descriptors, including SIFT, LIOP, DAISY, RI, and GB, are considered. Table 2 summarizes the performances in mAP of all the approaches on the four benchmarks. The first four approaches to image matching, i.e., SM (Leordeanu and Hebert, 2005), ACC (Cho et al., 2009), HV (Chen et al., 2013), VFC (Ma et al., 2014), consider a single descriptor at a time. Their performances with each of the five descriptors as well as the average performances are reported. The other four baselines and our approach can jointly take multiple descriptors into account, so only the performances of descriptor fusion are reported. In Table 2, the best performance on each benchmark is given in bold, while the best performance by using a single descriptor is given in italic and comes with a star sign.

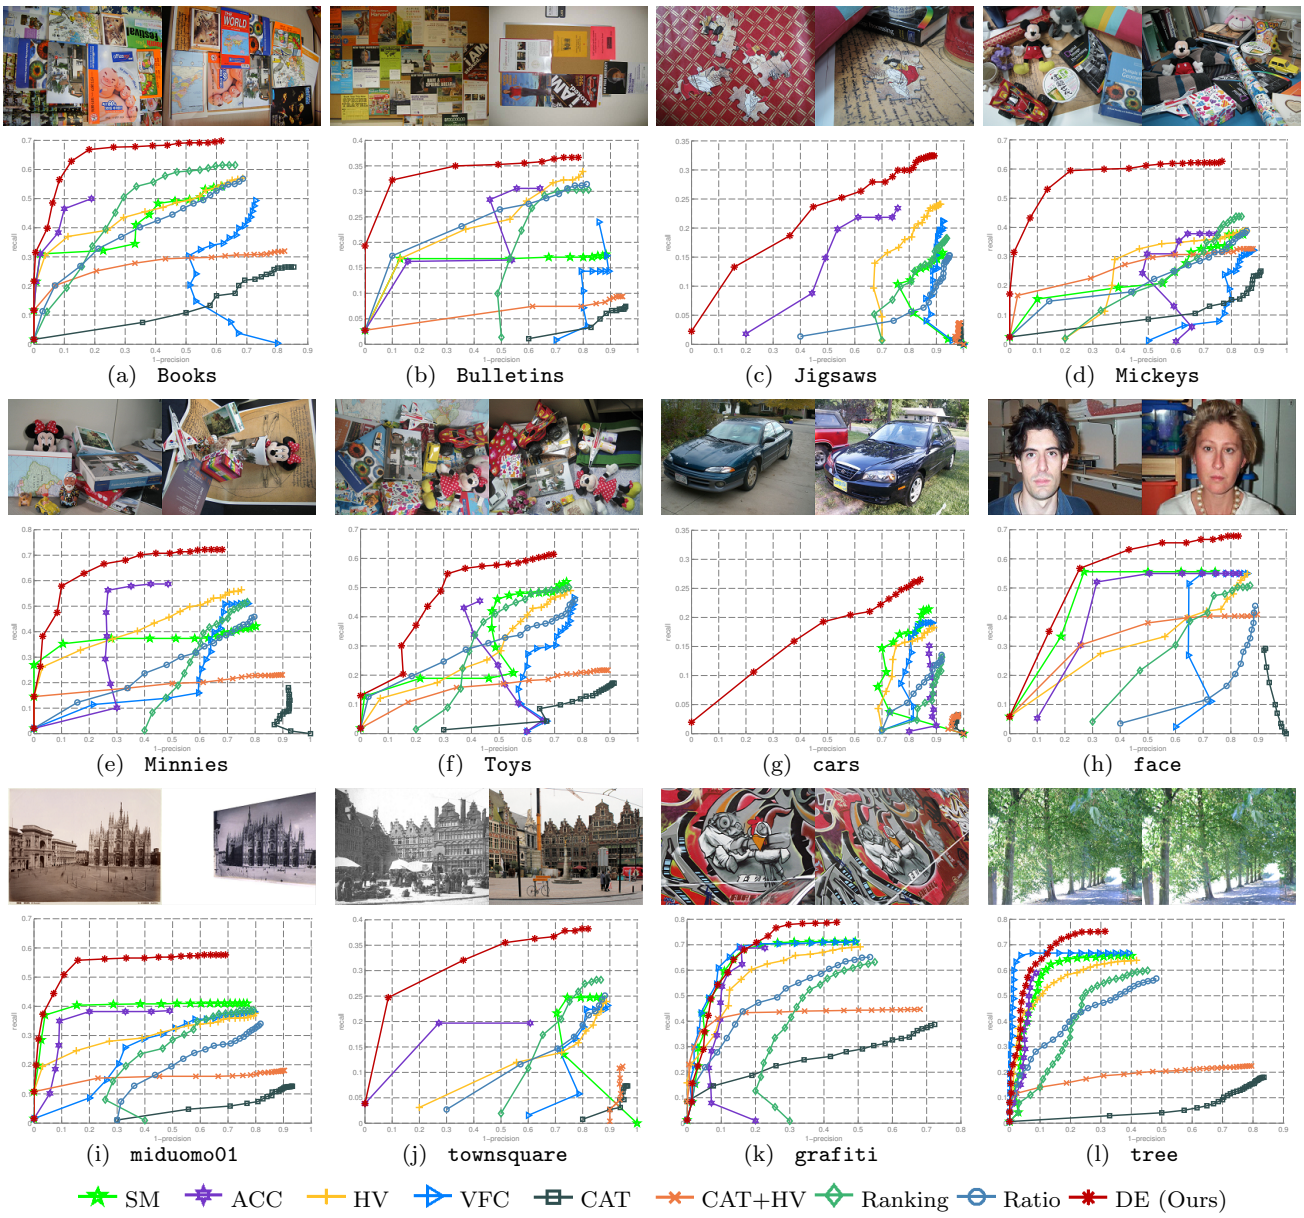

Fig. 5 The precision-recall curves on 12 image pairs. (a) ∼ (f) Six image pairs of the Co-reg dataset. (g) and (h) Two image pairs of the Object dataset. (i) and (j) Two image pairs of the SYM dataset. (k) and (l) Two image pairs of the VGG dataset.

We firstly focus on the cases where a single descriptor is used. SIFT gives the best performance in dataset Co-reg, while LIOP performs best in dataset VGG. The experimental results show that the five descriptors complement each other and no single descriptor can get the best performance on all the four datasets. The optimal descriptor for matching vary from image to image. It hence points out that fusing multiple descriptors can be a feasible way for improving performance. As for the performances of the image matching algorithms, baseline HV averagely gets the superior results, since it fully supports multiple object matching, and stably works with various descriptors.

The four baselines for descriptor fusion, i.e., CAT, CAT+HV, Ranking and Ratio, perform diversely. Baselines CAT and CAT+HV give poor performance. Even their accuracies in mAP fall behind those by the first four image matching algorithms that work with a single descriptor. It reveals that concatenation is not a good strategy for descriptor fusion, because worse descriptors degrade the discriminative power of the concatenated descriptor. Baselines Ranking and Ratio, especially Ratio, lead to much better matching results.

Table 2 The performances in mAP of the eight baselines and our approach on the four datasets

method	descriptor	dataset			
		Co-reg	Object	SYM	VGG
SM	SIFT	47.59	21.90	16.31	60.93
	LIOP	21.16	14.61	15.96	70.06
	DAISY	35.76	20.91	17.85	71.43
	RI	14.61	10.37	15.35	67.60
	GB	9.75	14.18	16.91	65.57
	Average	25.77	16.40	16.48	67.12
ACC	SIFT	*60.28**	21.81	*21.73**	77.96
	LIOP	29.83	7.41	19.27	81.92
	DAISY	36.49	17.88	20.85	81.63
	RI	15.10	7.51	18.21	77.29
	GB	8.88	16.16	19.53	74.02
	Average	30.12	14.15	19.92	78.56
HV	SIFT	60.12	23.73	18.73	77.95
	LIOP	43.97	15.44	17.28	82.41
	DAISY	50.06	22.32	18.84	82.74
	RI	37.14	14.55	16.78	78.69
	GB	12.08	*30.24**	19.18	76.40
	Average	40.67	21.25	18.16	79.64
VFC	SIFT	31.11	21.44	18.77	79.67
	LIOP	11.79	9.26	15.25	*84.24**
	DAISY	16.29	18.74	16.37	83.40
	RI	4.51	4.46	13.40	74.54
	GB	1.77	21.60	15.42	72.19
	Average	13.09	15.10	15.84	78.81
CAT	All	19.62	6.20	7.22	58.76
CAT+HV	All	39.70	15.75	11.90	70.62
Ranking	All	48.48	19.38	23.00	82.05
Ratio	All	53.61	18.62	23.62	85.07
DE (Ours)	All	**78.46**	**40.88**	**37.31**	**88.39**

The two baselines averagely outperform the four image matching algorithms, but still fall behind them if the best descriptor in each dataset is chosen. For instance, baseline Ratio gives 53.61% in Co-reg and 18.62% in Object, while baseline ACC with SIFT achieves 60.28% in Co-reg and baseline HV with GB achieves 30.24% in Object. In contrast, our approach allows mutual verification across different descriptors in an unsupervised manner, and correct correspondences will distinguish themselves with high coherence to each other. The quantitative results show that our approach can make the most of fusing various feature descriptors, and achieve significant performance gains over all the baselines on the four datasets.

The mAP summarizes the performances of matching approaches on the whole dataset. To look inside how they work on individual images, precision-recall curves (precisely 1−precision vs. recall curves here) are used. The Co-reg dataset consists of six image pairs. The resulting precision-recall curves by all the approaches are shown in Fig. 5(a) ∼ 5(f). Note that we manually pick the best descriptor for baselines SM, ACC, HV and VFC to draw their curves for the sake of clearness. Thus, their performances may be overestimated in this sense. Baselines ACC and HV can deal with multiple object matching, while baseline VFC and SM are less robust in the cases. Ranking and ratio can increase the recall with the aid of multiple descriptors in most cases, but their precision is unsatisfactory. Our method can effectively match multiple objects, and considerably boost both the recall and precision by leveraging multiple descriptors.

We also select two pairs of images from each of the other three datasets, and plot the corresponding precision-recall curves in Fig. 5(g) ∼ 5(l), respectively. The three datasets have different types of variations in matching, such as intra-class variations in dataset Object, combined changes in dataset SYM, and imaging condition changes in dataset VGG. Our approach can deal with these variations by adaptively picking appropriate descriptors in matching interest points, and result in the superior performance. The exceptions are shown in Fig. 5(k) `grafiti` and 5(l) `tree`, two examples from dataset VGG. As can be seen in Table 2, descriptor LIOP achieves satisfactory results, and dominates the other descriptors on dataset VGG. Hence, the performance gain of our approach is not significant in the cases.

7.3 Visualization of Matching Results

To gain insight into the quantitative results, we show the matching results by our approach and the adopted baselines on a few images. Fig. 6 displays the matching results as well as the recalls by our approach, on image pair `Books` of the Co-reg dataset, with each of the five adopted descriptors and all of them. The correct correspondences are colored with their colors indicating the descriptors by which they are established, i.e., SIFT in orange, LIOP in blue, DAISY in green, RI in magenta, and GB in cyan. Descriptor SIFT finds the most correct matchings in this example. As shown in Fig. 6(f), our approach indeed selects most correspondences established by SIFT. Similar observation can be found in Fig. 7, in which the matching results on image pair `dragonfly` of the Object dataset are plotted. In this example, descriptor GB performs best, and our

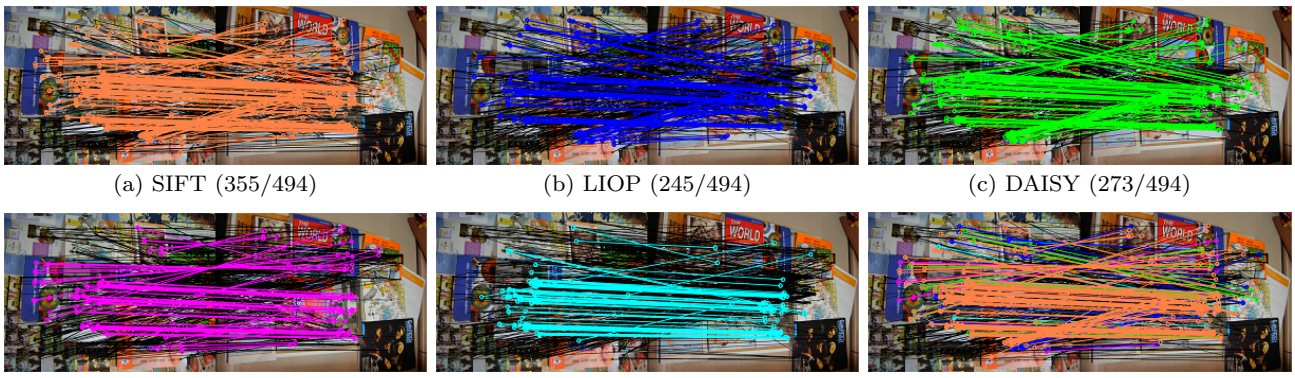

(a) SIFT (355/494) (b) LIOP (245/494) (c) DAISY (273/494)

(d) RI (175/494) (e) GB (145/494) (f) All (419/494)

Fig. 6 The matching results by our approach, on image `Books` of the Co-reg dataset, with (a) the SIFT descriptor, (b) the LIOP descriptor, (c) the DAISY descriptor, (d) the RI descriptor, (e) the GB descriptor, and (f) all the five descriptors. The recalls are shown in brackets. The correct correspondences are colored, and their colors indicate the descriptors by which they are established.

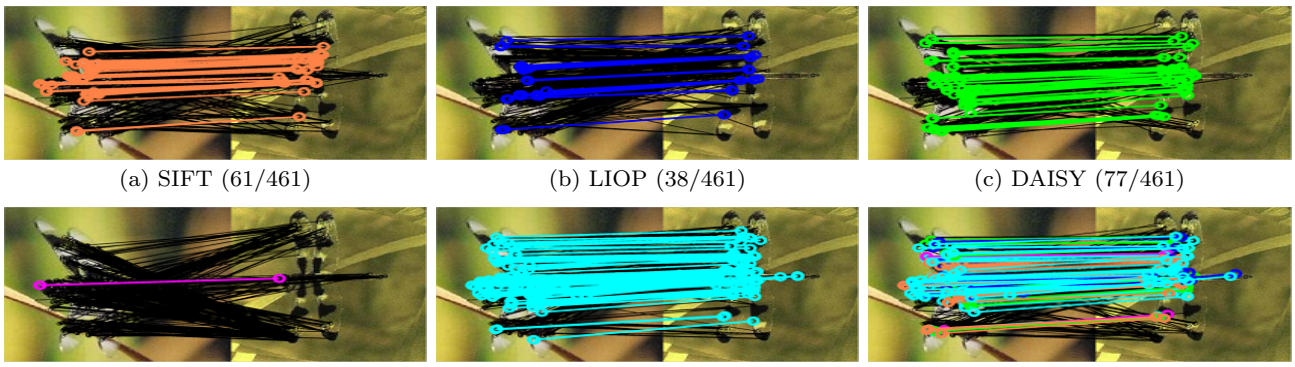

(a) SIFT (61/461) (b) LIOP (38/461) (c) DAISY (77/461)

(d) RI (3/461) (e) GB (132/461) (f) All (168/461)

Fig. 7 The matching results by our approach, on image `dragonfly` of the Object dataset, with (a) the SIFT descriptor, (b) the LIOP descriptor, (c) the DAISY descriptor, (d) the RI descriptor, (e) the GB descriptor, and (f) all the five descriptors. The recalls are shown in brackets. The correct correspondences are colored, and their colors indicate the descriptors by which they are established.

approach also finds most correspondences by GB. The results in Fig. 6 and Fig. 7 show why our approach can leverage multiple, complementary descriptors to boost the matching performance: It adaptively determines the correct correspondences established by the better descriptors.

In Fig. 8, the matching results by our approach (DE) and the four image matching algorithms on two image pairs, `paintedladies12` and `trevi02`, of the SYM dataset are shown. Our approach yields more dense and accurate matchings (red correspondences), and outperforms the four matching algorithms even if their respective best descriptors on this dataset have been manually chosen. In Fig. 9, our approach and the four baselines for descriptor fusion are compared on two image pairs, `wall` and `grafiti`, of the VGG dataset. Our approach in both cases carries out geometric verification, and effectively reduces the numbers of false positives (black correspondences) yielded by individual descriptors. Therefore, it achieves more satisfactory results.

To summarize, the visualization of the matching results demonstrates that our approach can effectively leverage multiple descriptors: On the one hand, it allows geometric layout verification across heterogeneous descriptors, and hence results in higher precision. On the other hand, it increases the number of correct correspondence candidates with the aid of complementary descriptors, and leads to higher recall. These properties enable our approach to alleviate the unfavorable issues in image matching, such as multiple objects with dramatic perspective variations in the Co-reg dataset, large intra-class variations in the Object dataset, the combined changes of lighting conditions and rendering styles in the SYM dataset, and imaging condition changes in the VGG dataset.

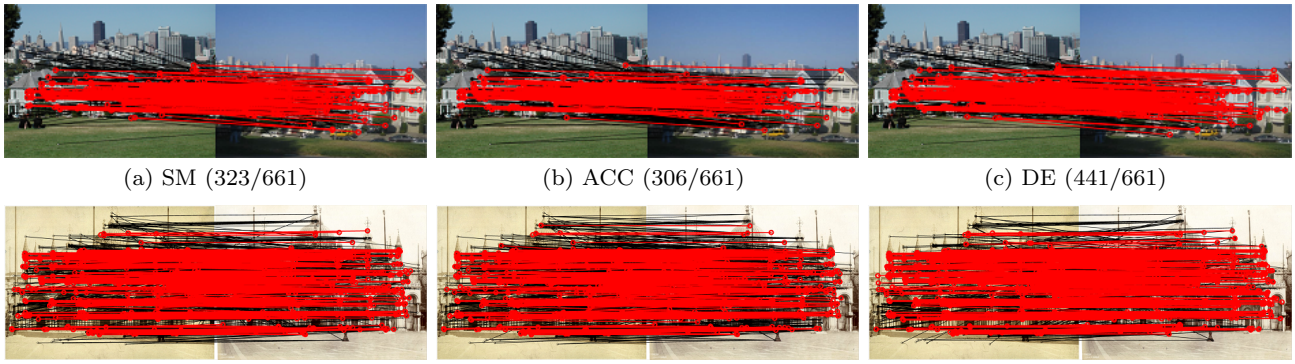

(a) SM (323/661) (b) ACC (306/661) (c) DE (441/661)

(d) HV(652/1546) (e) VFC (616/1546) (f) DE (859/1546)

Fig. 8 The matching results by our approach and the four image matching algorithms on two image pairs of the SYM dataset, including (a) ∼ (c) image pair `paintedladies12` and (d) ∼ (f) image pair `trevi02`.

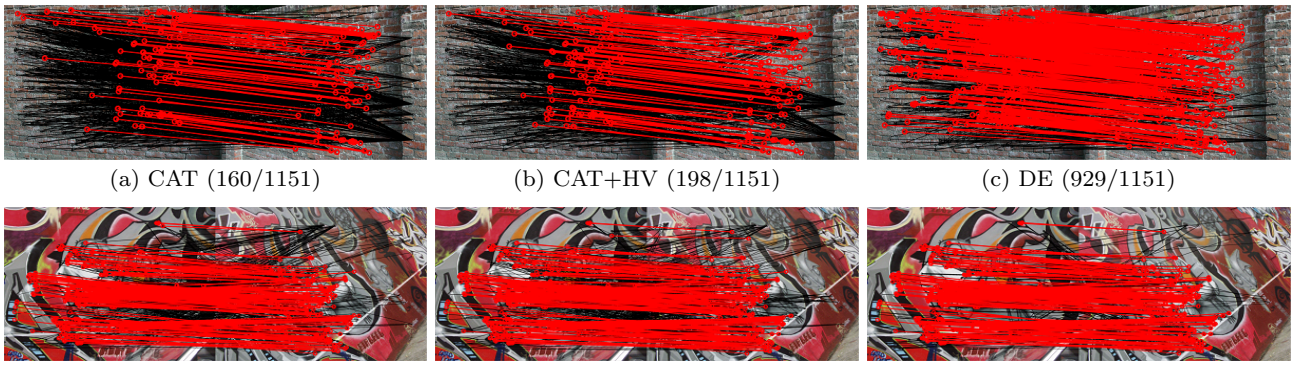

(a) CAT (160/1151) (b) CAT+HV (198/1151) (c) DE (929/1151)

(d) Ranking (521/821) (e) Ratio (537/821) (f) DE (649/821)

Fig. 9 The matching results by our approach and the four baselines for descriptor fusion on two image pairs of the VGG dataset, including (a) ∼ (c) image pair `wall` and (d) ∼ (f) image pair `grafiti`.

8 Conclusion

We have presented an effective approach that can leverage multiple, complementary descriptors, and boost the performance of image feature matching. Specifically, the correspondences yielded by all descriptors are firstly projected into the homography space, in which both geometric and spatial consistency among the correspondences are measured by computing the geodesic distance on a designed graph. One-class SVM is then employed to rank the correspondences according to their consensus with each other. The proposed approach is featured with high flexibility in the sense that it can work with any elliptical region detectors as well as heterogeneous descriptors. Besides, it selects plausible correspondences across descriptors in a fully unsupervised way: no prior knowledge about images to be matched is required. Our approach has been comprehensively evaluated on four benchmark datasets, and compared with both the state-of-the-art approaches to feature matching as well as the baselines for descriptor fusion. The experimental results demonstrate that our approach can significantly boost the matching quality in both precision and recall, and is superior to the existing approaches and the baselines. In the future, we plan to generalize and apply this work to computer vision and image processing applications where accurate and dense matchings are appreciated, such as image alignment, object recognition, and motion estimation.

References

Abdel-Hakim, A. E., Farag, A. A. (2006) CSIFT: A SIFT descriptor with color invariant characteristics. In Proc. Conf. Computer Vision and Pattern Recognition

Avrithis, Y., Tolias, G. (2014) Hough pyramid matching: Speeded-up geometry re-ranking for large scale image retrieval. Int. J. Computer Vision 107(1):1–19

Bay, H., Tuytelaars, T., Van Gool, L. (2006) SURF: Speeded up robust features. In Proc. Euro. Conf. Computer Vision

Belongie, S., Malik, J., Puzicha, J. (2002) Shape matching and object recognition using shape contexts.

IEEE Trans. on Pattern Analysis and Machine Intelligence 24(4):509–522

Berg, A., Malik, J. (2001) Geometric blur for template matching. In Proc. Conf. Computer Vision and Pattern Recognition

Bosch, A., Zisserman, A., Muñoz, X. (2007) Representing shape with a spatial pyramid kernel. In Proc. ACM Conf. Image and Video Retrieval

Brox, T., Malik, J. (2011) Large displacement optical flow: Descriptor matching in variational motion estimation. IEEE Trans. on Pattern Analysis and Machine Intelligence 33(3):500–513

Carneiro, G. (2010) The automatic design of feature spaces for local image descriptors using an ensemble of non-linear feature extractors. In Proc. Conf. Computer Vision and Pattern Recognition

Chang, C.-C., Lin, C.-J. (2001) LIBSVM: A library for support vector machines. ACM Trans. on Intelligent Systems and Technology 2(3):27:1–27:27, software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm

Chen, H.-Y., Lin, Y.-Y., Chen, B.-Y. (2013) Robust feature matching with alternate Hough and inverted Hough transforms. In Proc. Conf. Computer Vision and Pattern Recognition

Cho, M., Shin, Y.-M., Lee, K.-M. (2008) Co-recognition of image pairs by data-driven monte carlo image exploration. In Proc. Euro. Conf. Computer Vision

Cho, M., Lee, J., Lee, K.-M. (2009) Feature correspondence and deformable object matching via agglomerative correspondence clustering. In Proc. Int' Conf. Computer Vision

Cho, M., Lee, J., Lee, K.-M. (2010) Reweighted random walks for graph matching. In Proc. Euro. Conf. Computer Vision

Choi, O., Kweon, I.S. (2009) Robust feature point matching by preserving local geometric consistency. Computer Vision and Image Understanding 113(6):726–742

Chui, H., Rangarajan, A. (2003) A new point matching algorithm for non-rigid registration. Computer Vision and Image Understanding 89(2-3):114–141

Cour, T., Srinivasan, P., Shi, J. (2006) Balanced graph matching. In Proc. Advances in Neural Information Processing Systems

Cox, T. F., Cox, M. A. A. (1994) Multidimentional Scaling. Chapman & Hall, London

Datta, R., Joshi, D., Li, J., Wang, J. Z. (2008) Image retrieval: Ideals, influences, and trends of the new age. ACM Comput. Surv. 40(2)

Fischler, M., Bolles, R. C. (1981) Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM 24(6):381–395

Harada, T., Nakayama, H., Kuniyoshi, Y. (2010) Improving local descriptors by embedding global and local spatial information. In Proc. Euro. Conf. Computer Vision

Hauagge, D., Snavely, N. (2012) Image matching using local symmetry features. In Proc. Conf. Computer Vision and Pattern Recognition

Hua, G., Brown, M., Winder, S. (2007) Discriminant embedding for local image descriptors. In Proc. Int' Conf. Computer Vision

Kadir, T., Zisserman, A., Brady, M. (2004) An affine invariant salient region detector. In Proc. Euro. Conf. Computer Vision

Leordeanu, M., Hebert, M. (2005) A spectral technique for correspondence problems using pairwise constraints. In Proc. Int' Conf. Computer Vision

Leordeanu, M., Sukthankar, R., Hebert, M. (2012) Unsupervised learning for graph matching. Int. J. Computer Vision 96(1):28–45

Li, H., Huang, X., He, L. (2013) Object matching using a locally affine invariant and linear programming techniques. IEEE Trans. on Pattern Analysis and Machine Intelligence 35(2):411–424

Liu, H., Yan, S. (2010) Common visual pattern discovery via spatially coherent correspondences. In Proc. Conf. Computer Vision and Pattern Recognition

Lobaton, E., Vasudevan, R., Alterovitz, R., Bajcsy, R. (2011) Robust topological features for deformation invariant image matching. In Proc. Int' Conf. Computer Vision

Lowe, D. (2004) Distinctive image features from scale-invariant keypoints. Int. J. Computer Vision 60(2):91–110

Ma, J., Zhao, J., Tian, J., Yuille, A., Tu, Z. (2014) Robust point matching via vector field consensus. IEEE Trans. on Image Processing 23(4):1706–1721

Mikolajczyk, K., Schmid, C. (2004) Scale and affine invariant interest point detectors. Int. J. Computer Vision 60(1):63–86

Mikolajczyk, K., Schmid, C. (2005) A performance evaluation of local descriptors. IEEE Trans. on Pattern Analysis and Machine Intelligence 27(10):1615–1630

Mikolajczyk, K., Tuytelaars, T., Schmid, C., Zisserman, A., Matas, J., Schaffalitzky, F., Kadir, T., Van Gool, L. (2005) A comparison of affine region detectors. Int. J. Computer Vision 65(1-2):43–72

Mortensen, E., Deng, H., Shapiro, L. (2005) A SIFT descriptor with global context. In Proc. Conf. Computer Vision and Pattern Recognition

Pang, S., Xue, J., Tian, Q., Zheng, N. (2014) Exploiting local linear geometric structure for identifying correct mtches. Computer Vision and Image Understanding

128(0):51–64
Schmid, C., Mohr, R. (1997) Local grayvalue invariants for image retrieval. IEEE Trans. on Pattern Analysis and Machine Intelligence 19(5):530–534
Schölkopf, B., Platt, J., Shawe-Taylor, J., Smola, A., Williamson, R. (2001) Estimating the support of a high-dimensional distribution. Neural Computation 13(7):1443–1471
Shechtman, E., Irani, M. (2007) Matching local self-similarities across images and videos. In Proc. Conf. Computer Vision and Pattern Recognition
Szeliski, R., Shum, H.-Y. (1997) Creating full view panoramic image mosaics and environment maps. In ACM Trans. on Graphics (Proc. SIGGRAPH)
Tenenbaum, J., de Silva, V., Langford, J. (2000) A global geometric framework for nonlinear dimensionality reduction. Science 290(5500):2319–2323
Tola, E., Lepetit, V., Fua, P. (2010) DAISY: An efficient dense descriptor applied to wide-baseline stereo. IEEE Trans. on Pattern Analysis and Machine Intelligence 32(5):815–830
Varma, M., Ray, D. (2007) Learning the discriminative power-invariance trade-off. In Proc. Int' Conf. Computer Vision
Wang, Z., Fan, B., Wu, F. (2011) Local intensity order pattern for feature description. In Proc. Int' Conf. Computer Vision
Weinzaepfel, P., Revaud, J., Harchaoui, Z., Schmid, C. (2013) Deepflow: Large displacement optical flow with deep matching. In Proc. Int' Conf. Computer Vision
Winder, S., Brown, M. (2007) Learning local image descriptors. In Proc. Conf. Computer Vision and Pattern Recognition
Zhang, J., Marszalek, M., Lazebnik, S., Schmid, C. (2007) Local features and kernels for classification of texture and object categories: A comprehensive study. Int. J. Computer Vision 73(2):213–238
Zhang, W., Wang, X., Zhao, D., Tang, X. (2012) Graph degree linkage: Agglomerative clustering on a directed graph. In Proc. Euro. Conf. Computer Vision
Zheng, Y., Doermann, D. (2006) Robust point matching for nonrigid shapes by preserving local neighborhood structures. IEEE Trans. on Pattern Analysis and Machine Intelligence 28(4):643–649